

SIZE-CHANGE TERMINATION, MONOTONICITY CONSTRAINTS AND RANKING FUNCTIONS

AMIR M. BEN-AMRAM

School of Computer Science, The Tel-Aviv Academic College, Israel
e-mail address: benamram.amir@gmail.com

REVISION NOTE. This is a revised and corrected version of the article originally published on July 11, 2010.

ABSTRACT. Size-Change Termination (SCT) is a method of proving program termination based on the impossibility of infinite descent. To this end we may use a program abstraction in which transitions are described by *monotonicity constraints* over (abstract) variables. Size-change graphs are a subclass where only constraints of the form $x > y'$ and $x \geq y'$ are allowed. Both theory and practice are now more evolved in this restricted framework than in the general framework of monotonicity constraints. This paper shows that it is possible to adapt and extend some theory from the domain of size-change graphs to the general case, thus complementing previous work on monotonicity constraints. In particular, we present precise decision procedures for termination; and we provide a procedure to construct explicit global ranking functions from monotonicity constraints in singly-exponential time, which is better than what has been published so far even for size-change graphs.

1. INTRODUCTION

This paper is concerned with termination analysis. This is a fundamental and much-studied problem of software verification, certification and transformation. While being, historically, the epitome of undecidability, much progress has been made on automatizing termination proofs, so that now we find termination assertions in practical programming tools such as JML [24] and proof assistants such as ACL2 [31] and Isabelle [23]. One of the contributing factors to the development of automated termination analysis was its importance in designing certain meta-programs, e.g., interpreters for Logic Programs [32, 16] and partial evaluators [20, 21]. In such applications, the danger is that the meta-program will diverge, which is unacceptable. In theorem provers, termination proofs are required for showing that recursive definitions are sound.

A subproblem of termination analysis is the construction of *global ranking functions*. Such a function is required to decrease in each step of a program (for “step” read basic block, function call, etc, as appropriate); the function witnesses the progress towards termination.

1998 ACM Subject Classification: D.2.4; F.3.1.

Key words and phrases: program analysis, SCT, termination, ranking functions.

An explicitly presented ranking function whose descent is (relatively) easy to verify is a useful *certificate* for termination [22] and may have other uses, such as running-time analysis [2].

As the general halting problem is undecidable, every method for termination analysis consists, in principle, of identifying some *subproblem* that is decidable. This can be done in a more or less structured manner; a very structured approach is to break the termination problem for programs into two stages: the first is *abstraction* of the program, in which the concrete program is replaced with an abstract one, essentially a simplified model of the original. The second stage is *analysis* of the abstract program. One benefit of this separation is that the abstract programs may be rather independent of the concrete programming language. Another one is that the termination problem for the abstract programs may be decidable.

Size-change termination (SCT [26]) is such an approach. It views a program as a transition system with states. The abstraction consists in forming a *control-flow graph* for the program, identifying a set of *state variables*, and forming a finite set of *size-change graphs* that are abstractions of the transitions of the program. In essence, a size-change graph is a set of inequalities between variables of the source state and the target state. Thus, the SCT abstraction is an example of a transition system defined by *constraints* of a particular type.

The technique concentrates on well-founded domains, where infinite descent is impossible. An SCT termination proof is a proof that any (hypothetical) infinite run would decrease some value monotonically and endlessly so that well-foundedness is contradicted. Since only chains of non-increasing values are of interest, only two types of inequalities were admitted into the constraints in [26]: $x > y'$ (old value of x greater than new value of y) and $x \geq y'$.

Size-change graphs lend themselves to a very natural generalization: Monotonicity Constraints. Here, a transition may be described by any conjunction of order relations, including equalities as well as strict and non-strict inequalities, and involving any pair of variables from the source state and target state. Thus, it can express a relation among source variables, that applies to states in which the transition may be taken; a relation among the target variables, which applies to states which the transition may produce; and, as in SCT, relations involving a source variable and a target variable, but here equalities can be used, as well as relations like $x \leq x'$, that is, an increase.

The *Monotonicity Constraint Systems* treated in this paper will include another convenience, *state invariants* associated with a point in the control-flow graph. These too are conjunctions of order constraints.

Monotonicity constraint systems generalize the SCT abstraction and are clearly more expressive. It may happen that analysis of a program yields monotonicity constraints which are not size-change graphs; in such a case, simply approximating the constraints by a size-change graph may end up missing the termination proof. For an example, see the next section. It is not surprising, perhaps, that Monotonicity Constraints actually predated the SCT framework—consider the Prolog termination analyses in [27, 13]. But as often happens in science, concentrating on a simplified system that is sufficiently interesting was conducive to research, and the formulation of the SCT framework led to a series of interesting discoveries. To pick up some of the salient points:

- The SCT abstraction has a simple semantics, in terms of transition systems.

- It has a useful combinatorial representation as a set of graphs, known as the size-change graphs.
- A termination criterion has been formulated in terms of these graphs (the existence of an infinite descending thread in every infinite multipath [26]).
- This criterion is equivalent to the termination of every model (transition system)—in logical terms, this condition is sound and complete [26, 25].
- Termination of a set of size-change graphs can be effectively decided; while having exponential worst-case time, the method is often usable. The best-known method is a closure-based algorithm [26], which can be theoretically related to termination proofs by disjunctively well-founded transition invariants [33, 12]. More theoretically, the complexity of the problem has been identified as PSPACE-complete [26].
- It has been established that a global ranking function can also be effectively constructed from the size-change graphs. Lee [25] gave the first proof, where the size of the resulting ranking expression is up to triply-exponential in the size of the abstract program. This left open the challenging problem of improving this upper bound. Progress regarding certain special cases is reported in [8].

Which of the useful results can also be obtained in the stronger framework of monotonicity constraints? One contribution of this paper is an answer: *in essence, all of them*.

A second contribution of this paper is an algorithm to verify termination of a monotonicity constraint system while constructing a global ranking function, all in singly-exponential time. Thus, we solve the open problem from [25], and, surprisingly, by tackling a super-problem¹. The ranking functions generated have a simple form, and are based solely on lexicographic descent. Their worst-case size as well as ordinal height match known lower bounds.

Prior Work. The Prolog termination analysers *Termilog* [27, 28] and *Terminweb* [13] can both construct monotonicity constraints for a program. The abstraction used in both systems also includes instantiation patterns, particular to logic programming, however ignoring this aspect, the termination analysis boils down to deciding termination of monotonicity constraint systems.

The implementation of SCT analysis in ACL2 [31] added to the original SCT formulation a notion of state invariants (in the form of “calling contexts”). This predates our inclusion of state invariants in the monotonicity-constraint abstraction.

Codish, Lagoon and Stuckey [12] were the first to investigate the question of monotonicity constraints versus size-change graphs. They made the intriguing observation, that the termination test used in [27, 13, 28] is sound and complete for SCT, but incomplete for general monotonicity constraints. They also presented a correct test (that is, sound and complete), closely related to the algorithms presented here (see Section 5.3).

Usefulness of monotonicity constraints. Another important point in [12] is how monotonicity constraints can imply termination for the integer domain, in a way that transcends the well-founded model discussed above. For example, monotonicity constraints are sufficient for deducing termination of a loop such as `while(x < y) x = x + 1`. The theory presented in this paper can, in principle, be adapted to the integer domain. The basic ideas can be found

¹The author is aware of another singly-exponential solution for SCT, communicated privately by Chin Soon Lee.

in [12] and [5]. The details are, however, complicated, and in order to keep this paper lucid (and of reasonable length), they are deferred to a future publication. Since this paper is an attempt to present a coherent theory for monotonicity constraints, including several novel definitions and results, it will be restricted to the well-founded model, which allows for a more elegant theoretical development. It may be true that for proving termination in the cases where the well-founded model is natural (programs operating over lists, trees etc.), SCT works so well that the extension is practically redundant (at least, I have not seen empirical evidence to the contrary). Nonetheless, users of SCT in such contexts may benefit of the new ranking-function construction, as no singly-exponential worst-case construction has been known for SCT; and future application to the integer domain will be facilitated by the theoretical foundation.

The reader is asked to bear, therefore, with the rather abstract presentation in this paper, and turn to [26] or other papers on SCT for concrete examples of programs and their abstract representation. All examples proved to terminate using size-change graphs will also be amenable to the new ranking-function construction.

Some of the results in this paper have been presented at CAV 2009 [5].

How to abstract a program. Let us conclude this discussion by saying something about the method of abstracting programs, that is, extracting the monotonicity constraints. Such an analysis may be done in a very simple-minded manner in the style of the toy example given in [26], but when it is necessary to establish relations between *computed values*, a stronger static analysis will be necessary. Fortunately, such tools exist. A classic example is [10] for logic programs, which underlies Termilog. The more recent [9] is used in Terminweb, as is polyhedral analysis [15] which is an all-times favourite for handling imperative programs [4, 1, 36].

The implementations of SCT analysis in ACL2 [31] and Isabelle [22] rely on a theorem prover to derive constraints.

In the rest of this paper, we concentrate on analysing the abstract programs, and do not aim to contribute to the art of abstraction, except indirectly, by drawing attention to the possibilities of abstracting a program to a monotonicity constraint system.

2. BASIC DEFINITIONS AND EXAMPLES

This section introduces monotonicity constraint systems (MCS) and their semantics, and formally relates them to the SCT abstraction. As terminology is not uniform across SCT-related publications, some arbitrary choices had to be made. For instance, we shall use the term *flow point* where some references use *flow-chart point*, *program location* or *function*, the latter obviously in a functional-programming context.

Throughout the text, the symbol \triangleright will serve as a meta-variable ranging over relations $\{>, \geq\}$.

2.1. Monotonicity constraint systems: structure. A monotonicity constraint system is an abstract program. An abstract program is, essentially, a set of *abstract transitions*. An abstract transition is a relation on (abstract) program states.

When describing program transitions, it is customary to mark the variables in the *resulting* state with primes (e.g., x'). For simplicity, we will name the variables x_1, \dots, x_n (regardless of what flow point we are referring to). Of course, in practice there is no reason

for the number of variables to be the same throughout the program, but this does not affect the theory in any essential way.

Definition 2.1. A *monotonicity constraint system*, or MCS, is an abstract program representation that consists of a control-flow graph (CFG), monotonicity constraints and state invariants, all defined below.

- A control-flow graph is a directed graph (allowing parallel arcs) over the set F of *flow points*.
- A *monotonicity constraint* (MC) is a conjunction of order constraints $x \triangleright y$ where $x, y \in \{x_1, \dots, x_n, x'_1, \dots, x'_n\}$.
- Every CFG arc $f \rightarrow g$ is associated with a monotonicity constraint G . We write $G : f \rightarrow g$.
- For each $f \in F$, there is an *invariant* I_f , which is a conjunction of order constraints among the variables $\{x_1, \dots, x_n\}$.

In writing order constraints, we will also use $<, \leq$ as a syntactic sugar ($x < y$ is $y > x$), and $x = y$ to mean $x \geq y \wedge y \geq x$.

The terms “abstract program”, “constraint system” and “MCS instance” are used interchangeably, when context permits. The letters \mathcal{A}, \mathcal{B} are usually used to denote such a program; $F^{\mathcal{A}}, F^{\mathcal{B}}$ denote their respective flow-point sets.

We next show how an MC is represented by a labeled digraph (directed graph); the notation $x \xrightarrow{r} y$ represents an arc from x to y with label r .

Definition 2.2 (constraints as graphs). The graph representation of a monotonicity constraint is a labeled digraph (V, E) with $V = \{x_1, \dots, x_n, x'_1, \dots, x'_n\}$ and E includes a labeled arc for each constraint: specifically, for a constraint $x \triangleright y$, an arc $x \xrightarrow{\triangleright} y$.

The labeled arcs are referred to, verbally, as strict (for $>$) and non-strict (\geq).

Note that arcs may connect two source variables, two target variables or a source and a target variable—in any direction. The notation $x \rightarrow y$ may be used to represent an arc from x to y (of unspecified label). In diagrams, to avoid clutter, we distinguish the types of arcs by using a dashed arrow for the weak inequalities or equalities (see Figure 1).

An equality constraint $x = y$ is represented by a pair of non-strict arcs. In certain algorithms, it is convenient to assume that the such arcs are distinguished from “ordinary” non-strict arcs. We refer to them as *no-change arcs*.

Henceforth, we identify a MC with its graph representation: it is a graph and a logical conjunction of constraints at the same time. Context will usually clarify what view is taken,.

Example 2.3. Consider the following (contrived) program fragment:

```
while (x,y,z > 0)
  if (y>x)
    y = z;   x = unknown();  z = x-1
  else
    z = z-1; x = unknown();  y=x-1
```

The program computes over non-negative integers, which justifies the well-founded model. For representing this program as an MCS in an economic way, we transform each basic block (namely each branch of the *if*) into an MC. The CFG thus consists of a single




Figure 1: MCS example. The CFG (not shown) consists of a single flow-point and two self-loops. In the MC graphs, the left-hand side is the source. Broken arcs represent non-strict descent.

flow-point, representing the top of the loop, with two self-loops, G_1 and G_2 :

$$G_1 : \quad x < y \wedge z = y' \wedge x' > z'$$

$$G_2 : \quad x \geq y \wedge z > z' \wedge x' > y'$$

Figure 1 shows the graph representation of this MCS. □

One may wonder why only conjunctions are allowed. In fact, any Boolean combination of constraints can be put into disjunctive normal form, and then split into several MCs, one for each disjunct. Of course, there are combinations (e.g., in conjunctive normal form) that will be blown up exponentially. Apparently, this has not a problem in applications of SCT so far.

Another natural question is why state invariants are included in the definition, as it is possible to include the constraints I_f in every MC that transitions from f , making the state invariant redundant. However, it may be convenient (and is very much so in the algorithms of Sections 5–6) to make the association of certain constraints with a flow point, rather than a specific transition, explicit. Note also that static analysis algorithms (e.g., interval analysis [15]) often associate invariants with locations in the program, taking into account all transitions that lead to a given location.

The size-change graphs of [26] are a simple class of monotonicity constraints:

Definition 2.4 (size-change graph). A size-change graph (SCG) is a monotonicity constraint consisting only of relations of the forms $x_i \triangleright x'_j$. As a graph, it is bipartite and includes only arcs from source variables to target variables.

An *SCT instance* is a MCS where all constraints are size-change graphs and all invariants are trivial.

2.2. Semantics and termination. Recall that a well-order is a total order with no infinite strictly-descending chains, in other words, a well-founded total order. The semantics of the abstract program assumes a well-ordered set Val as the domain for variables' values. The non-strict order relation on Val is \geq and its strict part is denoted by $>$ (that is, $x > y \iff x \geq y \wedge x \neq y$).

While it seems that most applications of SCT use total orders (e.g., comparing data objects by their size), there are exceptions (notably in conjunction with term rewriting systems). In fact, all notions and results in this paper work equally well with partial orders,

and even partial quasi-orders (where \geq is not antisymmetric), except for the global ranking functions, which presume a total order. For uniformity, I chose to make total order the basic assumption.

Definition 2.5 (states). Let \mathcal{A} be an n -variable MCS. A *state* of \mathcal{A} (or an abstract state) is a pair (f, σ) , where $f \in F^{\mathcal{A}}$ and $\sigma : \{1, \dots, n\} \rightarrow \text{Val}$ represents an assignment of values to the variables.

Satisfaction of a predicate e with free variables x_1, \dots, x_n (for example, $x_1 > x_2$) by an assignment σ is defined in the natural way, and expressed by $\sigma \models e$. If e is a predicate involving the $2n$ variables $x_1, \dots, x_n, x'_1, \dots, x'_n$, we write $\sigma, \sigma' \models e$ when e is satisfied by setting the unprimed variables according to σ and the primed ones according to σ' .

Definition 2.6 (transitions). A transition is a pair of states, a *source state* s and a *target state* s' . For $G : f \rightarrow g \in \mathcal{A}$, we write $(f, \sigma) \xrightarrow{G} (g, \sigma')$ if $\sigma \models I_f$, $\sigma' \models I_g$ and $\sigma, \sigma' \models G$. We say that transition $(f, \sigma) \mapsto (g, \sigma')$ is *described by* G .

G is called *unsatisfiable* if it describes no transition.

The *transition system* associated with \mathcal{A} is the relation $T_{\mathcal{A}}$ defined by

$$(s, s') \in T_{\mathcal{A}} \iff s \xrightarrow{G} s' \text{ for some } G \in \mathcal{A}.$$

Practically, the transition system would be an abstraction of the transitions of a concrete program. The variables of flow points might represent actual data in the program, but quite often they are already an abstraction, like the size of a concrete object (this is, in fact, the source for the name of the SCT method).

Definition 2.7 (run). A *run* of $\mathcal{T}_{\mathcal{A}}$ is a (finite or infinite) sequence of states $\tilde{s} = s_0, s_1, s_2 \dots$ such that for all i , $(s_i, s_{i+1}) \in \mathcal{T}_{\mathcal{A}}$.

Note that by the definition of $\mathcal{T}_{\mathcal{A}}$, a run is associated with a sequence of CFG arcs labeled by G_1, G_2, \dots where $s_{i-1} \xrightarrow{G_i} s_i$. This sequence constitutes a walk in the CFG (recall that a walk is a directed path where repeated nodes and arcs are allowed).

Definition 2.8 (termination). Transition system $\mathcal{T}_{\mathcal{A}}$ is *uniformly terminating* if it has no infinite run.

MCS \mathcal{A} is said to be *terminating* if $\mathcal{T}_{\mathcal{A}}$ is uniformly terminating for any choice of Val .

Definition 2.9 (\vdash). Let $P(s, s')$ be any predicate over states s, s' , possibly written using variable names, e.g., $x_1 > x_2 \wedge x_2 < x'_2$. We write $G \vdash P$ if $\forall s, s' : s \xrightarrow{G} s' \Rightarrow P(s, s')$.

Definition 2.10. A *global ranking function* for a transition system \mathcal{T} with state space St is a function $\rho : St \rightarrow W$, where W is a well-founded set, such that $\rho(s) > \rho(s')$ for every $(s, s') \in \mathcal{T}$.

A ranking function for a MCS \mathcal{A} is a ranking function for $T_{\mathcal{A}}$. Namely, it satisfies $G \vdash \rho(s) > \rho(s')$ for every $G \in \mathcal{A}$.

The qualifier *global* may be omitted in the sequel, since this paper does not deal with the notion of local ranking functions.




Figure 2: A multipath.

Example 2.3. (continued). The MCS shown on Page 6 is terminating. To see this, note first that the set of constraints $s \xrightarrow{G_2} s' \wedge s' \xrightarrow{G_1} s''$ is unsatisfiable (the reader is invited to verify this by writing the constraints down, or by inspecting the graphs). Thus, a G_2 -transition will never be followed by a G_1 ; in other words, a valid run can switch from G_1 to G_2 but cannot switch back. In G_2 , variable z decreases; it also decreases in G_1 , *if followed by G_1 again* (the reader is, again, invited to verify this).

We conclude that the following function descends (lexicographically) in every possible transition, and constitutes a ranking function for this program:

$$\rho(\mathbf{x}, \mathbf{y}, \mathbf{z}) = \begin{cases} \langle 1, \mathbf{z} \rangle & \text{if } \mathbf{y} > \mathbf{x} \\ \langle 0, \mathbf{z} \rangle & \text{if } \mathbf{y} \leq \mathbf{x}. \end{cases}$$

Note that the MCs are not size-change graphs. Their best approximations as size-change graphs are $G_1^{\text{scg}} = \{z > y'\}$ and $G_2^{\text{scg}} = \{z > z'\}$, which do not prove termination. In fact, the issue of unsatisfiable combinations of constraints (as for G_2 followed by G_1) never arises with SCT instances; this is one of the salient differences between the MC abstraction and the SCT one.

3. MULTIPATHS, WALKS AND TERMINATION

The purpose of this section is to formulate the *combinatorial* termination condition for monotonicity constraint systems, that is, the property that MC graphs should satisfy so that termination of any associated transition system can be deduced. This approach was very useful in the study of the SCT abstraction, and we would like to define the corresponding notions for MCS so that for the special case of SCT, they match the known results.

The size-change termination principle [26] states that a program is known to terminate if “in every (hypothetic) infinite run, something descends infinitely”. In order to analyse such an infinite run, we analyse sequences of size-change graphs (and more generally, MC graphs) that describe such runs. For this purpose, we introduce the concept of a *multipath* (following [26]).

3.1. The Criterion for Monotonicity Constraint Systems.

Definition 3.1 (multipath). Let \mathcal{A} be an n -variable MCS, and let $f_0 \xrightarrow{G_1} f_1 \xrightarrow{G_2} f_2 \dots$ be an MC-labeled path in the CFG. The *multipath* M that corresponds to this path is a (finite or infinite) graph with nodes $x[t, i]$, where t ranges from 0 up to the length of the path, and $0 < i \leq n$. Its arcs are obtained by merging the following sets: for all $t \geq 1$, M includes the arcs of G_t , with source variable x_i renamed to $x[t-1, i]$ and target variable x'_j renamed to

$x[t, j]$. In addition, for all $t \geq 0$, arcs representing the invariant I_{f_t} are included (with node $x[t, i]$ representing x_i).

The multipath may be written concisely as $G_1 G_2 \dots$; if M_1, M_2 are finite multipaths, M_1 corresponding to a CFG path that ends where M_2 begins, we denote by $M_1 M_2$ the result of concatenating them in the obvious way.

Figure 2 depicts multipath $G_1 G_2 G_1$, based on the MCs from Figure 1.

Clearly, a multipath represents a conjunction of constraints on a set of variables associated with its nodes. We can consider assignments σ to these variables, where the value assigned to $x[t, i]$ may be denoted $\sigma[t, i]$; such an assignment may satisfy the constraints, or not. A satisfying assignment defines a concrete run of \mathcal{T}_A , along the given CFG path.

If we start at any node in a multipath, and walk along arcs, we are tracing a descending chain of values. This way, walks in the multipath may be used to prove when infinite multipaths are unsatisfiable.

Definition 3.2. A walk that includes a strict arc is said to be *descending*. A walk that includes infinitely many strict arcs is *infinitely descending*.

Definition 3.3 (termination criterion). An MCS \mathcal{A} is *size-change terminating* if every infinite \mathcal{A} -multipath has an infinitely descending walk.

Note that the walk above may actually be a cycle! In this case it is contained in a finite section of the multipath, and, logically, it implies the condition $x > x$ for some variable x . Thus, such a multipath is unsatisfiable and corresponds to no concrete run. If the walk is not a cycle, it indicates an infinite descending chain of values and this contradicts well-foundedness. Once again, we conclude that the multipath is unsatisfiable. We conclude that if \mathcal{A} is size-change terminating, it can have no infinite runs.

Lemma 3.4. *If MCS \mathcal{A} is size-change terminating, it is (semantically) terminating.*

Proof. Suppose that \mathcal{A} is size-change terminating. For any (hypothetical) infinite run \tilde{s} of T_A there is an underlying infinite path in the CFG, which induces a multipath M ; the values assigned to variables in \tilde{s} should satisfy all constraints expressed by M , which is impossible because of the infinitely descending walk. □

Thus, size-change termination is a sound criterion for termination.

To show completeness, we suppose that an infinite multipath without infinite descent exist, and show the existence of a domain *Val* over which this infinite multipath is satisfiable. This is achieved using the following definition and lemma.

Definition 3.5. Let S be any set. A binary relation \succ on S is a *strict order* if it is transitive and irreflexive. A binary relation \succeq is a *non-strict quasi-order* if it is transitive and reflexive; if it is also antisymmetric, it becomes a (partial) order—the word *partial* tacitly applies to all the above.

Let \succ be a strict order and \succeq a non-strict quasi-order. The relations are *compatible* if

$$a \succ b \Rightarrow a \succeq b \tag{3.1}$$

$$a \succ b \wedge b \succeq c \Rightarrow a \succ c \tag{3.2}$$

$$a \succeq b \wedge b \succ c \Rightarrow a \succ c \tag{3.3}$$

It seems to be “folklore” that a well-founded partial order can be extended to a well-founded total order (given the axiom of choice or some similar machinery). The next lemma makes a slightly stronger statement (involving a quasi-order and a compatible strict order), and is given with proof².

Lemma 3.6. *Let S be any set, \succeq a quasi partial order on S , and \succ a compatible well-founded strict order. There is a well-ordered set B (with order \geq and strict $>$) and a mapping $h : S \rightarrow B$ that agrees with the order relations, that is:*

$$x \succeq y \Rightarrow h(x) \geq h(y)$$

$$x \succ y \Rightarrow h(x) > h(y)$$

Proof. Let $X = \mathcal{P}(S) - \emptyset$. For each $x \in X$, let M_x be the set of minimal elements of x under \succ . This is always non-empty. For $x \in X$ and $a \in x$, let $U_a^x = \{b \in x \mid b \preceq a\}$; note that if $a \in M_x$ then $U_a^x \subseteq M_x$ and $a \in U_a^x$, so U_a^x is not empty.

Using the Axiom of Choice, let f be a choice function with $f(x) \in M_x$ for all x . Define, by transfinite induction over *the class of all ordinals*, the partial function with range $\mathcal{P}(S) \setminus \{\emptyset\}$:

$$g(\beta) = U_{f(y)}^y, \quad \text{where} \quad y = S - \bigcup_{\gamma < \beta} g(\gamma),$$

unless $y = \emptyset$ or $g(\gamma)$ is undefined for some $\gamma < \beta$.

Define, by the ZF *axiom of replacement*, $B = \{\beta \mid \exists x \in S : g(\beta) = x\}$. Since B is a set of ordinals, it cannot contain all ordinals (by the Burali-Forti paradox), thus there is an ordinal α not in B . Identifying an ordinal with the set of smaller ordinals, $(\alpha + 1) \setminus B$ is a non-empty set of ordinals. Since the ordinals are well ordered, there is a least ordinal $\beta \in (\alpha + 1) \setminus B$; in fact it is the least ordinal not in B . Therefore $g(\beta)$ is undefined. It cannot be that the second “unless” clause in the definition holds (since β is the least such ordinal), so it must be that $S - \bigcup_{\gamma < \beta} g(\gamma) = \emptyset$, and therefore for every $a \in S$ there is some $\gamma < \beta$ such that $a \in g(\gamma)$. In fact, by the definition of B , such γ is unique, so letting $h(a) = \gamma$ defines a total function $h : S \rightarrow B$.

We claim that h agrees with the order relations. To see this, let $a, b \in S$. By the definition of h ,

$$a \in g(h(a)) \subseteq M_y \subseteq y = S - \bigcup_{\gamma < h(a)} g(\gamma).$$

Suppose that $a \succ b$. If $h(a) \leq h(b)$, then also $b \in y$, so a is not minimal in y and is not in M_y , a contradiction. Thus, $h(a) > h(b)$, so $>$ agrees with \succ . Next, suppose that $a \succeq b$; if $b \in y$, then (by the definitions of M_y and $U_{f(y)}^y$), we have $b \in U_{f(y)}^y = g(h(a))$; so $h(b) = h(a)$. If $b \notin y$, then $h(b) < h(a)$. Either way, \geq is seen to agree with \succeq . \square

Lemma 3.7. *MCS \mathcal{A} is terminating only if it is size-change terminating.*

Proof. Suppose that \mathcal{A} is not size-change terminating. Hence, an infinite multipath M can be formed, without infinite descent. Let S be the set of nodes of M , identified by the usual notation $x[t, i]$. Our aim is to show that there is a model for M , that is, M is satisfiable.

² The proof is inspired by a proof of the well-ordering theorem in <http://planetmath.org/encyclopedia/ProofOfZermelosWellOrderingTheorem.html> (dated 2007-03-12).

We first define order relations on S . Specifically, define the relation \succeq on S by: $x[t, i] \succeq x[s, j]$ if and only if M includes a walk from $x[t, i]$ to $x[s, j]$. Define the relation \succ on S by: $x[t, i] \succ x[s, j]$ if and only if M includes a descending walk from $x[t, i]$ to $x[s, j]$.

It is easy to verify that (S, \succeq, \succ) satisfy the assumptions of the last lemma. We choose Val to be the well-ordered set B from the conclusion of the lemma, and σ to be the mapping h ; then σ is an assignment that satisfies M . \square

Combining Lemmas 3.4 and 3.7, we have

Theorem 3.8. *MCS \mathcal{A} is terminating if and only if it is size-change terminating.*

3.2. The SCT criterion. The SCT condition [26] is similar to the MCs termination condition (Definition 3.3), but only concerns walks that proceed forward in the multipath. Obviously, with SCT graphs, there are no other walks anyway.

Definition 3.9. In a multipath, a thread is a walk that only includes arcs in a forward direction ($x_i \rightarrow x'_j$). We say that MCS \mathcal{A} *satisfies SCT* if every infinite \mathcal{A} -multipath has an infinitely descending thread.

Note that for a general MCS, there is a difference between being size-change terminating and “satisfying SCT.” The latter refers to the criterion defined above.

Example 2.3 shows that the SCT condition, while clearly a sufficient condition for termination, is not a necessary one when general monotonicity constraints are considered.

4. DECIDING TERMINATION: THE CLOSURE ALGORITHM

We will show in this section an algorithm to decide MCS termination. This algorithm is very similar to algorithms used in previous works such as [34, 27, 17] which, however, are not complete (see Section 4.5).

4.1. Consequence-closure and composition.

Definition 4.1. A monotonicity constraint G is *closed under logical consequence* (or just closed) if, whenever $G \vdash x > y$, for $x, y \in \{x_1, \dots, x_n, x'_1, \dots, x'_n\}$, this constraint is explicitly included in G ; and if $G \vdash x \geq y$, but not $x > y$, then $x \geq y$ is included in G .

Note that for $G : f \rightarrow g$, the condition $G \vdash P$ takes the invariants I_f and I_g into account (consider Definitions 2.6 and 2.9). Thus, a closed MC subsumes the invariants in its source and target states.

Definition 4.2. An MC H is *at least as strong as* G if whenever $G \vdash P$, also $H \vdash P$. The (consequence) *closure* of a monotonicity constraint G , denoted \overline{G} , is the weakest MC that is at least as strong as G and is consequence-closed.

Practically, calculating the closure means inserting into G all the relations that can be deduced from G . This is easy given the graph representation of the MC, as described next.

Definition 4.3. The *weighted-graph representation* of an MC $G : f \rightarrow g$ consists of the same nodes as G , and an arc $x \rightarrow y$ for every relation $x \triangleright y$ included in G , I_f or I_g ³. Each non-strict arc is given a weight of 0 and each strict arc, a weight of -1 .

³Parallel arcs can be eliminated, preferring $>$ to \geq .

Lemma 4.4. *Let G^{wgt} be the weighted-graph representation of MC G . Then G is unsatisfiable if and only if a negative-weight cycle exists in G^{wgt} . Assuming that G is satisfiable, \overline{G} includes an arc $x \rightarrow y$ if and only if a path from x to y exists in G^{wgt} ; this arc is strict if there is a path of negative weight.*

We omit the straight-forward justification of this lemma. It implies that a standard All-Pairs Lightest-Path (weighted shortest-path) algorithm, such as the Floyd-Warshall algorithm [14], can be used to find if G is satisfiable, and compute \overline{G} if it is, in polynomial time—specifically, $O(n^3)$.

Remark: this is not best possible, asymptotically, but is a simple solution. Different improvements may be tried, but whether they are of practical value cannot be judged on a theoretical level, and would only make sense if the graphs are large. For example, one may divide the running time by a constant, dependent on the machine word length, using a bit-level representation. For sparse graphs with $a \ll n$ arcs, one may opt for an $O(na)$ solution based on repeated DFS search (details are left to the reader).

Definition 4.5 (composition). The *composition* of MC $G_1 : f \rightarrow g$ with $G_2 : g \rightarrow h$, written $G_1; G_2$, is a MC with source f and target h , which includes all the constraints among s, s' implied by $\exists s'' : s \xrightarrow{G_1} s'' \wedge s'' \xrightarrow{G_2} s'$.

Composition, too, can be implemented efficiently by a lightest-path algorithm applied to the multipath $G_1 G_2$. As already noted, the procedure also determines whether the multipath is satisfiable (note that, as Example 2.3 demonstrates, two satisfiable MCs may form an unsatisfiable multipath).

The same graph procedure can be applied to any finite multipath $M = G_1 \dots G_\ell$. It computes $\overline{M} = \overline{G_1}; \dots; \overline{G_\ell}$, which we call the *collapse* of M (if $\ell = 1$, it is the consequence-closure of G_1). Thus, \overline{M} includes an arc $x_i \rightarrow x'_j$ if and only if M includes a path from node $x[0, i]$ to node $x[\ell, j]$; and similarly for the other combinations $(x_i \rightarrow x_j, x'_i \rightarrow x'_j$ and $x'_i \rightarrow x_j)$. The arc is strict if and only if the path includes a strict arc. This gives us the following observation

Observation 4.6. Consider an infinite \mathcal{A} -multipath M , represented as the concatenation of finite segments $M_1 M_2 \dots$, and let $M' = (\overline{M}_1)(\overline{M}_2) \dots$. If M' has an infinite descending walk, so does M .

Definition 4.7 (closure set). Given an MCS \mathcal{A} , its closure set \mathcal{A}^* is

$$\{\overline{M} \mid M \text{ is a satisfiable } \mathcal{A}\text{-multipath}\}.$$

The set \mathcal{A}^* is finite, since there are finitely many possible MCs.

4.2. A termination test.

Definition 4.8 (cyclic multipath). We say that a multipath M (possibly a single MC) is *cyclic* if its source and target flow-points are equal. This is equivalent to stating that MM is a valid multipath.

Definition 4.9 (circular variant). For an MC G , the *circular variant* G° of G is a directed graph obtained by adding, for every parameter x_i , an edge $x_i \leftrightarrow x'_i$. This edge is treated as a pair of no-change arcs, but is distinguished from any arcs already present in G . These additional edges are called *shortcut edges*.

This definition is meant to be used with a cyclic MC. The shortcut edges are used to analyze the effect of juxtaposing multiple copies of G (as in the multipath GG). This will be clearer in the proof of Theorem 4.12.

Definition 4.10 (types of cycles). Let G be a cyclic MC. A cycle in G° is a path commencing and ending at the same node. It is a *forward cycle* if it traverses shortcut edges more often in the backward direction (from x'_i to x_i) than it does in the forward direction⁴.

A *balanced cycle* is a cycle that traverses shortcut edges equally often in both directions.

A cycle is *descending* if it includes a strict arc.

Definition 4.11 (Local Termination Test). We say that G passes the *Local Termination Test*, or LTT, if G has a descending cycle, either forward or a balanced.

See Figure 3, Parts (a)-(c), for an example of a forward descending cycle. If arc $x'_3 \xrightarrow{\geq} x'_1$ were replaced with $x'_3 \xrightarrow{\geq} x'_2$, a balanced cycle would ensue.

Using the graph representation, the Local Termination Test is not hard to implement; let us sketch the algorithm, leaving the proof and fine details to the interested reader. First, we note that one can test one strongly connected component (SCC) of G° at a time, and only if it includes a strict arc. Secondly, we claim that if there is any forward cycle in the SCC, then it contains a forward descending cycle. This case can be identified by assigning weights to arcs and looking for a negative-weight cycle using, say, the Bellman-Ford shortest-path algorithm [14]; this takes $O(na)$ time if the graph has n nodes and a arcs⁵. If no cycle of this kind exists, there can only be a balanced descending cycle, which would constitute a zero-weight cycle; having computed the single-source shortest-path distances, such cycles can easily be found since they are cycles in the shortest-path graph.

Theorem 4.12. *MCS \mathcal{A} is size-change terminating if and only if every cyclic MC in \mathcal{A}^* passes the Local Termination Test.*

Proof. For the forward implication (the “if”), suppose that every cyclic MC in \mathcal{A}^* passes the Local Termination Test. Consider an infinite \mathcal{A} -multipath $M = G_1 G_2 \dots$ and assume, by way of contradiction, that it is satisfiable.

Consider the set of positive integers, and label each pair (t, t') , where $t < t'$, by

$$G = \overline{G_t; G_{t+1}; \dots G_{t'-1}}$$

which must be in \mathcal{A}^* , since this multipath is satisfiable. Thus every pair has a label, and the set of labels is finite. By Ramsey’s theorem (in its infinite version), there is an infinite set of positive integers, I , such that all pairs (t, t') with $t, t' \in I$ carry the same label G_I .

Thus for any $t, t' \in I$ with $t < t'$, $\overline{G_t; G_{t+1}; \dots G_{t'-1}} = G_I$. By Observation 4.6, it now suffices to show that multipath $(G_I)^\omega$ (infinite sequence of G_I ’s) has an infinite descending walk.

Let $v_1, e_1, v_2, e_2, \dots, e_{s-1}, v_s$ be the nodes and arcs (alternatingly) of the descending cycle in G_I , where each v_j is either x_{i_j} or x'_{i_j} for some index i_j . We can map the cycle onto a walk in $(G_I)^\omega$, as follows. The first node is $x[s, i_1]$. If the arc e_1 is an ordinary arc of G , the walk follows this arc to $x[s+1, i_2]$, $x[s, i_2]$ or $x[s-1, i_2]$ (depending on the direction of the arc). If e_1 is a shortcut arc, the walk is not extended: v_2 is also mapped to $x[s, i_1]$

⁴this naming may seem strange, but will later see that such a cycle is “unwound” into a forward-going walk in the multipath G^ω .

⁵A better result, at least theoretically, is $O(an^{1/2})$ for n nodes and a arcs, due to Goldberg [19].




Figure 3: (a) An MC G , (b) its circular variant G^o ; (c) the directions of shortcut edges are set to form a forward descending cycle, (d) a walk in a prefix of G^ω , corresponding to the cycle. The notation x_i^t is a shorthand for $x[t, i]$. (Example after Codish, Lagoon and Stuckey)

(necessarily $i_2 = i_1$). We proceed in this manner until we complete the cycle and return to v_1 . At this point, our walk will have reached a node $x[s', i_1]$ for some s' . The relation among the forward and backward crossings of shortcut edges implies that $s' \geq s$; specifically $s' = s$ in the “balanced” case, and $s' > s$ in the “forward” case (hence the name). In the former case, we have discovered a descending cycle in $(G_I)^\omega$ which indicates unsatisfiability. In the latter, we have a walk from $x[s, i_1]$ to $x[s + d, i_1]$ for some positive d ; another such walk can be added to reach $x[s + 2d, i_1]$, and so on; we conclude that an infinitely-descending walk occurs in $(G_I)^\omega$. This establishes the forward implication (see Figure 3).

For the converse implication, suppose that MCS \mathcal{A} is size-change terminating, and let G be a cyclic MC in \mathcal{A}^* . Consider the multipath G^ω . Note that $G = \overline{M}$ for some finite, satisfiable \mathcal{A} -multipath M . Let ℓ be the length of M . Consider the infinite \mathcal{A} -multipath M^ω , obtained by concatenating infinitely many copies of M , to which we shall refer as *blocks*. By assumption, M^ω has an infinite descending walk. If this walk is a cycle, contained entirely within one of the blocks, then M is unsatisfiable, and \overline{M} will not appear in \mathcal{A}^* . Thus, the walk must cross block boundaries. Let us concentrate on the variables $x[t, i]$ occurring on these boundaries. Since the walk is infinitely descending, some index i has to occur twice on the walk, say at nodes $x[t\ell, i]$ and $x[t'\ell, i]$ where $t \leq t'$, and there has to be a strict arc between these two points⁶. There is thus a descending walk from $x[t\ell, i]$ to $x[t'\ell, i]$. Such a walk can be transformed into a cycle C in G^o in the following way: starting with $x[t\ell, i]$, pick the first segment S of the walk (at least a single arc) that ends up again on a block boundary. If S ends up in $x[t + \ell, j]$ include in C the G arc $x_i \rightarrow x'_j$ followed by a backward shortcut arc to x_j . If S ends up in $x[t, j]$ for some $j \neq i$ then C includes either an arc

⁶These considerations apply equally to walks that are cycles and to “open” walks.

$x_i \rightarrow x_j$ or an arc $x'_i \rightarrow x'_j$. In the former case, include $x_i \rightarrow x_j$ in C ; in the latter, include the path $x_i \rightarrow x'_i \rightarrow x'_j \rightarrow x_j$ where the first and last arcs are shortcuts. Finally, if S ends up in $x[t - \ell, j]$, include in C the path $x_i \rightarrow x'_i \rightarrow x_j$ where the first arc is a (backward) shortcut. This process continues until we get to $x[t'\ell, i]$ in the original walk, whereupon C will end at x_i , becoming a cycle. The inequality $t' \geq t$ implies that the number of backward shortcuts in C will not exceed the number of forward shortcuts.

We conclude that a descending cycle, either balanced or forward, exists in G° . \square

Algorithm 4.13. (Deciding termination of an MCS \mathcal{A})

- (1) Build \mathcal{A}^* by a transitive closure procedure:
 - (a) Initialize a set \mathcal{S} to $\{\overline{G} \mid G \in \mathcal{A} \text{ is satisfiable}\}$.
 - (b) For any $G : f \rightarrow g$ and $H : g \rightarrow h$ in \mathcal{S} , include also $G;H$ in \mathcal{S} , unless it is unsatisfiable.
 - (c) Repeat the above step is until no more elements can be added to \mathcal{S} .
 At this point, \mathcal{S} is \mathcal{A}^* (we omit a detailed proof).
- (2) For each cyclic G in \mathcal{S} , apply the Local Termination Test to G . Pronounce failure (non-termination) if a graph that fails the test is found.
- (3) If the previous step has completed, the MCS terminates.

4.3. Complexity of Algorithm 4.13. The determining factor in the worst-case complexity of the algorithm is the number of elements in \mathcal{A}^* . An easy upper bound on the number of possible MCs, ignoring the identity of the source and target flow-points, is 6^{2n^2} , as each MC has $2n$ variables and six possible relations among any pair of them ($x < y$, $x > y$, $x \geq y$, $x \leq y$, $x = y$ or none). Thus, the size of the closure set is at most $m^2 6^{2n^2}$, where m is the number of flow-points.

As long as the algorithm does not fail, any pair of elements that can be composed (i.e., MCs associated with CFG arcs $f \rightarrow g$ and $g \rightarrow h$, for any f, g, h) has to be considered (Step 1.2). To each such pair, we apply composition, and if the result is cyclic, the local termination test. Thus we obtain an upper bound of

$$O(m^3(6^{2n^2})^2 n^3) = O(m^3 6^{4n^2} n^3).$$

This is an over-estimate, however a significantly better bound (i.e., with $o(n^2)$ exponent) is not known for this algorithm. Note that the next section provides algorithms of better worst-case bounds: specifically, reducing the exponent from $\Theta(n^2)$ to $\Theta(n \log n)$.

We remark that the management of data structures is, in this case, not costly, since all we need is to maintain a set of MCs so that one can efficiently add an element while testing to see if it was not already there. This can be done, for instance, by a radix tree [14], where each operation takes linear time in the number of bits that describe an MC (that is, $O(n^2)$).

Despite the exponential time and space complexity, a similar algorithm for SCT has proved quite useful in practice; and there are techniques to improve its performance, that are also applicable here. Let us briefly discuss these (admitting that the ultimate test of such techniques is empirical).

An important technique is reducing the size of the set \mathcal{S} by subsumption [7, 18]. An MC G is said to subsume H if G is less constrained than H . That is, every transition described by H is also described by G . With consequence-closed graphs, testing subsumption is easy. And in this case, one can safely ignore H . There is a cost involved in finding out whether

subsumption occurs, whenever an element is to be added to \mathcal{S} . But the reduction in the size of \mathcal{S} seems to make this worthwhile.

Another sort of optimization appears in [17, 26]: instead of applying the LTT to every element of \mathcal{A}^* , it suffices to do so for the idempotent ones (elements G such that $G;G = G$). This strategy does not reduce the size of the closure set maintained, only the number of local tests. In SCT, it also reduces the complexity of the local test: as shown in [26], it suffices to look for in-situ descent, that is, an arc of the form $x_i \xrightarrow{>} x'_i$ (which yields a forward descending cycle of length 2). This is not the case for general monotonicity constraints; for a little example, note that the MC in Figure 3 is idempotent, but does not include in-situ descent.

A warning is in order: if the set \mathcal{S} is reduced by subsumption, it is incorrect to test just the idempotent graphs. That is, there are non-terminating instances where an idempotent counter-example will not be found, as subsumption will have removed it⁷. Thus one has to choose between applying subsumption or testing only idempotent graphs, and it seems clear that the former option has much greater impact.

4.4. Complexity of the decision problem. What is the complexity class of the decision problem: *MCS Termination*? Algorithm 4.13 takes exponential time and space, but a polynomial-space version is possible. It is most likely worst in practice, but for theoretical completeness, let us prove

Theorem 4.14. *The MCS Termination problem is PSPACE-complete.*

Proof. It is known [26] that the SCT Termination problem is PSPACE-hard, which also applies to MCS because SCT is a special case. To show that the problem is in PSPACE, we will outline a non-deterministic polynomial-space algorithm for the complement problem, that is, non-termination. The result will follow since (by Savitch’s theorem) $\text{coNPSpace} = \text{NPSpace} = \text{PSPACE}$.

Algorithm 4.13 can be seen as a search for a counter-example—a cycle in the CFG that induces a multipath that fails the test. The non-deterministic algorithm guesses such a cycle. In each step, it adds a transition to the cycle while computing the composition of the next MC with a graph that represents the collapse of the multipath traversed so far. Only this graph, along with the current flow-point, have to be maintained in memory. Whenever the current flow-point is the same as the initial one, the local termination test is applied. If at some point, an unsatisfiable MC results, the algorithm has failed to find a counter-example. Otherwise it continues until finding one. \square

Remark. Since both the MCS termination problem and SCT termination are PSPACE-complete, a polynomial-time reduction of the former to the latter is known to exist. However, it is not given explicitly, and it is not clear whether such a reduction can be implemented efficiently enough to be actually useful. On the other hand, the next section will show a reduction in *exponential time*, that, paradoxically, may be useful (in fact, the reduction produces an *easy instance* of SCT, one which can be analyzed in polynomial time, as shown in Section 6).

⁷Alexander Krauss, private communication.

4.5. Comparison to previous algorithms. As previously mentioned, there have already been termination analyzers that dealt with monotonicity constraints, but Codish et al. [12] pointed out that their decision procedures were not complete; they emphasized the case of integers (as was mentioned in the introduction), but the incompleteness already occurs in the well-founded model. In fact, those procedures are similar to Algorithm 4.13, and using the terminology in this section we can present them so that the change in the new algorithm becomes clear.

Definition 4.15. Let G be a cyclic MC. A cycle in G° is a *zig-zag cycle* if it alternates arcs of G with shortcut edges, the latter always in the backward direction (from x'_i to x_i).

Note that a zig-zag cycle alternates forward arcs of G ($x \rightarrow y'$) with shortcut edges. This does not mean, however, that other arcs in the MC are ignored, because they might give rise to forward arcs by transitivity.

Definition 4.16. We say that G passes *Sagiv's Test* if G has a zig-zag descending cycle.

Clearly, this test is a special case of the LTT, and it should be easy to see that they coincide for SCT graphs. The reader probably sees, already, why this test is incomplete. In fact, Figure 3 shows a terminating instance that it misses.

Sagiv's test is from [34] and was used in the Termilog system as described in [27], and the Terminweb system as described in [13]. Later, in [17], the algorithm was refined by observing that it suffices to test idempotent elements. This improves its coverage (consider the cyclic MC: $x > y \wedge x' < y'$) but it is still incomplete, by the same example.

It should be noted, that this comparison is based on a “translation” of the cited works to MCS terms. They actually analyze logic programs. As for the algorithm proposed in [12], the discussion is postponed to the next section where it fits more naturally.

5. FULLY ELABORATED SYSTEMS AND STABILITY

In this section we describe a procedure that while increasing the size of an abstract program, simplifies its termination proof; in fact, we get back to the SCT condition. This result is interesting theoretically, for understanding the relation between the two formalisms, and also forms a basis for the algorithm in the following section (that determines termination while constructing an explicit ranking function).

The procedure duplicates flow-points while refining their invariants; this means that computations that reach the same program location under different conditions (i.e., different ordering of the variable values) will be represented as reaching different flow-points in the abstract program. Similar transformations can be found in program analysis in various guises (for example, in [27] and subsequent works, flow points are queries to Prolog clauses and are duplicated for different instantiation patterns). When we apply the transformation to monotonicity constraint systems in a brute-force way, we obtain what we shall call a *fully elaborated system*. We will see, in this section and the next, that fully elaborated systems are MC systems of a particularly structured kind.

In order to express the correctness of a transformation on abstract programs we begin by defining “simulation.”

5.1. Simulation. We define the notions of simulation and bisimulation for transition systems of the kind used in this work. A transition system \mathcal{A} simulates another one \mathcal{B} if (informally speaking) they have the same transition sequences, up to the identity of flow points and the indexing of variables.

Definition 5.1. Let \mathcal{A}, \mathcal{B} be transition systems, with flow-point sets $F^{\mathcal{A}}, F^{\mathcal{B}}$ respectively, and both having states described by n variables over Val . We say that \mathcal{A} *simulates* \mathcal{B} if there is a relation $\phi \subseteq F^{\mathcal{B}} \times F^{\mathcal{A}}$ (“correspondence of flow-points”) and, for all $(f, g) \in \phi$, a bijection $\psi_{g,f} : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ (“variable renaming”) such that for every (finite or infinite) state-transition sequence $(f_1, \sigma_1) \mapsto (f_2, \sigma_2) \mapsto (f_3, \sigma_3) \mapsto \dots$ of \mathcal{B} there is a corresponding sequence $(g_1, \sigma'_1) \mapsto (g_2, \sigma'_2) \mapsto (g_3, \sigma'_3) \mapsto \dots$ of \mathcal{A} with $(f_i, g_i) \in \phi$ and $\sigma'_i = \sigma_i \circ (\psi_{g_i, f_i})$. We say that \mathcal{A} *bisimulates* \mathcal{B} if, in addition, for every (finite or infinite) state-transition sequence $(g_1, \sigma'_1) \mapsto (g_2, \sigma'_2) \mapsto (g_3, \sigma'_3) \mapsto \dots$ of \mathcal{A} there is a corresponding sequence $(f_1, \sigma_1) \mapsto (f_2, \sigma_2) \mapsto (f_3, \sigma_3) \mapsto \dots$ of \mathcal{B} , also with $(f_i, g_i) \in \phi$ and $\sigma'_i = \sigma_i \circ (\psi_{g_i, f_i})$.

Thus, \mathcal{A} bisimulates \mathcal{B} if they simulate each other via the same pair of mappings.

Definition 5.2. We say that an abstract program \mathcal{A} (bi-)simulates an abstract program \mathcal{B} if $\mathcal{T}_{\mathcal{A}}$ (bi-)simulates $\mathcal{T}_{\mathcal{B}}$, via mappings ϕ and ψ , as above.

We say that \mathcal{A} *simulates* \mathcal{B} *deterministically* if for every $f \in F^{\mathcal{B}}$ and assignment σ satisfying I_f there is a unique $g \in F^{\mathcal{A}}$ with $(f, g) \in \phi$ such that, letting $\sigma' = \sigma \circ (\psi_{g,f})$, assignment σ' satisfies I_g .

If \mathcal{A} bisimulates \mathcal{B} , and \mathcal{A} simulates \mathcal{B} deterministically, we say (for brevity) that \mathcal{A} *bisimulates* \mathcal{B} *deterministically*.

Determinism means that the invariants of different \mathcal{A} flow-points that simulate a given \mathcal{B} flow-point have to be mutually exclusive.

Observation 5.3. Suppose that \mathcal{S} bisimulates \mathcal{T} . Then \mathcal{S} uniformly terminates if and only if \mathcal{T} does.

5.2. Elaboration.

Definition 5.4 (full elaboration). An MCS \mathcal{A} is *fully elaborated* if the following conditions hold:

- (1) Each state invariant fully specifies the relations among all variables. That is, for $i, j \leq n$, one of $x_i = x_j$, $x_i < x_j$ or $x_i > x_j$ is implied by I_f .
- (2) Each MC is closed under logical consequence.
- (3) No MC in \mathcal{A} is unsatisfiable.

Since the state invariant fully determines the relations among all variables, we can *re-index* the variables into sorted order, so that the invariant becomes

$$x_1 \{ \leq \} x_2 \{ \leq \} \dots \{ \leq \} x_n. \quad (5.1)$$

Of course, the re-indexing has to be incorporated also in MCs incident to this flow-point, but this is straight-forward to do. Indexing the variables in sorted order has some convenient consequences, such as the having the property:

Definition 5.5. G has the *downward closure property* if for all $k < j$, $x_i \xrightarrow{\geq} x'_j \in G$ entails $x_i \xrightarrow{\geq} x'_k \in G$; and $x_i \xrightarrow{\geq} x'_j \in G$ entails $x_i \xrightarrow{\triangleright} x'_k \in G$ for some $\triangleright \in \{>, \geq\}$.

The number of possible orderings of n variables plays a role in the combinatorics of fully elaborated instances. Note that equalities are possible. Therefore, the number of orderings is not $n!$, but a slightly larger number called *the n th ordered Bell number* B_n . An easily proved upper bound is $B_n \leq 2n^{n-1}$ (consider two cases—no equalities, or at least one). For more details, in [35, Seq. A670]. We denote the set of these orderings by Bell_n , and assume that we fix some convenient representation so that “orderings” can be algorithmically manipulated (for example, a series of (in)equalities as in (5.1)).

The algorithm of full elaboration follows almost immediately from the definitions.

Algorithm 5.6. (full elaboration) Given an MCS \mathcal{B} , this algorithm produces a fully-elaborated MCS \mathcal{A} that bisimulates \mathcal{B} .

- (1) For every $f \in F^{\mathcal{B}}$, generate flow-points f_π where π ranges over Bell_n . Define the variable renaming function $\psi_{f_\pi, f}$ so that $\psi_{f_\pi, f}(i)$ is the i th variable in sorted order, according to π . Thus, I_{f_π} will have exactly the form (5.1).
- (2) Next, for every MC $G : f \rightarrow g$ in \mathcal{B} , and every pair f_π, g_ϖ , create a size-change graph $G_{\pi, \varpi} : f_\pi \rightarrow g_\varpi$ as follows:
 - (a) For every arc $x \rightarrow y \in G$, include the corresponding arc in $G_{\pi, \varpi}$, according to the variable renaming used in the two \mathcal{A} flow-points.
 - (b) Complete $G_{\pi, \varpi}$ by closure under consequences; unsatisfiable graphs (detected by the closure computation) are removed from the constructed system.

Example 5.7. Let the system \mathcal{B} consist of a single flow-point, say f , with $I_f = \text{true}$; and a single MC over two variables,

$$G : \quad x_1 > x'_1 \wedge x_2 \geq x'_2 \wedge x'_1 \geq x'_2.$$

In \mathcal{A} , we have $B_2 = 3$ flow-points f_π , namely: $f_{[x_1 < x_2]}$, $f_{[x_1 = x_2]}$, $f_{[x_1 > x_2]}$. For readability, let us denote the variables in \mathcal{A} by y_i instead of x_i : then y_i represents x_i in the first two flow-points, but in $f_{[x_1 > x_2]}$, the indices are exchanged, to obtain an increasing order of value.

Figure 4 shows three of the graphs $G_{\pi, \varpi}$, first when initially constructed—just copying the arcs from G according to the variable renaming, then after adding the invariants I_{f_π} and I_{f_ϖ} , and finally after closure under consequences. \square

Complexity. For an MCS \mathcal{B} , let $|\mathcal{B}|$ denote the number of abstract transitions (MCs) in \mathcal{B} (without loss of generality, $|\mathcal{B}| \geq |F^{\mathcal{B}}|$).

Lemma 5.8. *Any MCS \mathcal{B} with n variables at any point can be transformed into a fully-elaborated system \mathcal{A} , deterministically bisimulating \mathcal{B} , in $O(|\mathcal{B}|n^{2n+1})$ time and space.*

Proof. This follows from a straight-forward analysis of Algorithm 5.6. Each MC G yields $B_n^2 = O(n^{2n-2})$ offsprings $G_{\pi, \varpi}$, and the work invested in each is $O(n^3)$ (as explained in the last section). \square

	$G_{[x_1 < x_2], [x_1 = x_2]}$	$G_{[x_1 < x_2], [x_1 > x_2]}$	$G_{[x_1 = x_2], [x_1 > x_2]}$
Variables permuted	$\begin{array}{c} y_1 \longrightarrow y'_1 \\ \downarrow \\ y_2 \longrightarrow y'_2 \end{array}$	$\begin{array}{c} y_1 \longrightarrow y'_1 \\ \nearrow \downarrow \\ y_2 \longrightarrow y'_2 \end{array}$	$\begin{array}{c} y_1 \longrightarrow y'_1 \\ \nearrow \downarrow \\ y_2 \longrightarrow y'_2 \end{array}$
Invariants added	$\begin{array}{c} y_1 \longrightarrow y'_1 \\ \uparrow \\ y_2 \longrightarrow y'_2 \end{array}$	$\begin{array}{c} y_1 \longrightarrow y'_1 \\ \uparrow \\ y_2 \longrightarrow y'_2 \end{array}$	$\begin{array}{c} y_1 \longrightarrow y'_1 \\ \uparrow \\ y_2 \longrightarrow y'_2 \end{array}$
Closure	$\begin{array}{c} y_1 \longrightarrow y'_1 \\ \uparrow \\ y_2 \longrightarrow y'_2 \end{array}$	$\begin{array}{c} y_1 \longrightarrow y'_1 \\ \uparrow \\ y_2 \longrightarrow y'_2 \end{array}$	$\begin{array}{c} y_1 \longrightarrow y'_1 \\ \uparrow \\ y_2 \longrightarrow y'_2 \end{array}$

Figure 4: Full-elaboration example. Dashed arcs are non-strict.

5.3. Stability. The main result of this section is a proof that MCS termination can be reduced to SCT termination, and that fully-elaborating the system achieves this. We begin by formalizing the significant property, enjoyed by fully-elaborated systems, which allows for the simplified termination proof: this property is called *stability*.

Definition 5.9. An MCS \mathcal{A} is *stable* if (1) all MCs in \mathcal{A} are closed under logical consequence; (2) all MCs in \mathcal{A} are satisfiable; (3) For all $G : f \rightarrow g$ in \mathcal{A} , whenever $G \vdash x_i \triangleright x_j$ (a relation between source variables), that relation is included in I_f . (4) Similarly, if $G \vdash x'_i \triangleright x'_j$, that relation is included in I_g .

Note that just replacing every MC by its consequence-closure (if satisfiable) will satisfy conditions (1) and (2), but will not necessarily make the system stable, since it may fail to satisfy (3)-(4). In fact, these conditions may force flow-points to be duplicated, since two MCs coming out of f may disagree on the conditions that must be placed in I_f .

Observation 5.10. A fully elaborated system is stable.

Full elaboration can be seen as a brute-force way of “stabilizing” an MCS. A system can also be stabilized by an iterative fixed-point computation, which is likely to end up with less duplication of flow points and MCs. For completeness, such an algorithm is described in Section 5.5. But let us now first present the benefits of stability.

Lemma 5.11. *In a stable MCS, every finite multipath is satisfiable.*

Proof. Let \mathcal{A} be a stable MCS. Let M be a finite \mathcal{A} -multipath. For M to be unsatisfiable, it must include a descending cycle. We shall prove that if a descending cycle exists in M , then the *shortest* descending cycle must be contained in a single MC. But this would make the MC unsatisfiable, contradicting the assumption of stability.

Suppose, to the contrary, that the shortest descending cycle spans more than one MC. Suppose that it spans MCs G_a, G_{a+1}, \dots, G_b . We can assume that it includes a source node of G_a , for otherwise G_a is unnecessary (by stability, an arc among target nodes of G_a also appears in G_{a+1}). So, there is a node $\nu = x[a-1, i]$ on the cycle. The node that precedes




Figure 5: Three examples of walks from $x[t, i_{j_t}]$ to $x[t+1, i_{j_{t+1}}]$: in the first (topmost), t^* may be taken as either $t-1$, or $t+2$; in the second, only as $t+2$; in the third, as either t or $t+1$.

it on the cycle is a node of G_a , and so is the node that succeeds ν . Since all three nodes lie in the same MC, which is consequence-closed, this two-arc path can be replaced by a single arc, deleting ν from the cycle. Thus, the presumed shortest cycle is not shortest. \square

Theorem 5.12. *A stable MCS is terminating if and only if it satisfies SCT.*

Proof. Let \mathcal{A} be a stable MCS. If \mathcal{A} satisfies SCT, it is terminating, since the SCT condition is a special case of the MCS termination criterion. For the converse direction, assume that \mathcal{A} satisfies the MCS termination criterion; we will prove that it satisfies SCT. Let M be an infinite \mathcal{A} -multipath. We know that it has an infinitely descending walk. Lemma 5.11 shows that it cannot be a cycle. Therefore, it extends to infinity. We shall prove that, under this assumption, there is an infinitely descending thread.

The walk is made out of arcs $x[t_k, i_{j_k}] \rightarrow x[t_{k+1}, i_{j_{k+1}}]$ for $k = 0, 1, \dots$. For all $t \geq t_0$, let j_t be the first occurrence of t in the sequence; note that this is well defined. The walk is broken into segments leading from $x[t, i_{j_t}]$ to $x[t+1, i_{j_{t+1}}]$, of which infinitely many are descending. We claim that each of this segments can be replaced with a single arc, strict when appropriate. This implies that SCT is satisfied.

As there is a walk from $x[t, i_{j_t}]$ to $x[t+1, i_{j_{t+1}}]$, consider the shortest one (the shortest strict one, if appropriate). Suppose that it consists of more than one arc. Then there is a node $x[t^*, i_{j_{t^*}}]$ occurring inside the segment (not at its ends) such that the node preceding it, $x[t^-, i_{j_{t^-}}]$ and the node succeeding it, $x[t^+, i_{j_{t^+}}]$ satisfy either $t^-, t^+ \leq t^*$ or $t^-, t^+ \geq t^*$ (it may take a little reflection to see that such a node must exist; see Figure 5). Then consider the two arcs—from $x[t^-, i_{j_{t^-}}]$ to $x[t^*, i_{j_{t^*}}]$ and from $x[t^*, i_{j_{t^*}}]$ to $x[t^+, i_{j_{t^+}}]$; they must both lie in the same MC, and by consequence closure, there is a single arc (strict if any of the two arcs is) to replace these two arcs. Thus, the presumed shortest walk is not shortest.

We conclude that the shortest walk consists of a single arc, which is then a forward arc in the MC. We conclude that an infinitely descending *thread* exists in M , so that SCT is satisfied. \square

5.4. A decision algorithm. As an immediate corollary, we obtain a new algorithm to decide MCS termination. Namely,

Algorithm 5.13. (Deciding termination of an MCS \mathcal{A})

- (1) Stabilize the system (e.g., by full elaboration)
- (2) Apply an SCT decision algorithm.

Note that since we are deciding SCT, we can ignore any “backward” arcs ($x'_j \rightarrow x_i$), as well as the state invariants, in other words retain just SCT graphs. This observation may possibly be useful in optimizing an implementation.

Another natural expectation is that it would be desirable in practice to avoid full elaboration when possible, using a more economic stabilization procedure. Such a procedure is described next. We will discuss the efficiency of Algorithm 5.13 afterwards.

5.5. A more economic stabilization. Given an MCS \mathcal{B} , the following algorithm computes a stable system \mathcal{A} that bisimulates it. No variable renaming is used, so we only have to compute the set of flow-points and the mapping ϕ .

Algorithm 5.14. Stabilization by fixed-point computation

- (1) Initialize the system \mathcal{A} : flow-points in \mathcal{A} will be uniquely identified by a pair $\langle f, I \rangle$ where f is the corresponding \mathcal{B} flow-point and I is a state invariant. Initially, for every $f \in F^{\mathcal{B}}$, we have $\langle f, I_f \rangle$ in \mathcal{A} where I_f is the invariant from \mathcal{B} . Set ϕ to associate f with $\langle f, I_f \rangle$ and copy all MCs and invariants from \mathcal{B} .
- (2) Replace all MCs with their closure under logical consequence; if an MC is unsatisfiable, delete it.
- (3) Repeat the following process until instructed to stop:
 - (a) Search for an MC $G : \langle f, I_f \rangle \rightarrow \langle g, I_g \rangle$ which is not stable (checking for stability is a simple graph algorithm based on Lemma 4.4). If no such MC is found, stop.
 - (b) Since G is not stable, there is a relation, $x_i \triangleright x_j$, such that $G \vdash x_i \triangleright x_j$, but $x_i \triangleright x_j$ is not in I_f (or a similar situation on the g side).
 - (i) Suppose that the missing relation is $x_i > x_j$. Create two \mathcal{A} flow-points to replace $\langle f, I_f \rangle$: $\langle f, I_f \wedge (x_i > x_j) \rangle$ and $\langle f, I_f \wedge (x_i \leq x_j) \rangle$. If any of these points is already in \mathcal{A} , there is nothing more to do about it. For a point which is new, all MC's previously leading from and to $\langle f, I_f \rangle$ have to be copied to this new flow-point. Finally the old flow-point $\langle f, I_f \rangle$ is deleted.
 - (ii) If the missing relation is $x_i \geq x_j$, the flow-points created will be $\langle f, I_f \wedge (x_i \geq x_j) \rangle$ and $\langle f, I_f \wedge (x_i < x_j) \rangle$. The rest is as above.

It is not hard to see that the above loop will terminate. In fact, the number of \mathcal{A} flow-points corresponding to a given \mathcal{B} flow-point f is bounded (by B_n), so at some point the set of flow-points will stop growing, and then all size-change graphs must be stable.

Note that this algorithm will not change an SCT instance; SCT instances are stable.

5.6. Efficiency of analysing an MCS by reduction to SCT. To analyse the efficiency of Algorithm 5.13, we have to consider two questions. First, what is the cost of stabilization? The cost of full elaboration is $O(|\mathcal{B}|n^{2n+1})$, where \mathcal{B} is the input system. Algorithm 5.14 may take less, and it seems reasonable to expect it to take less in many instances (where there are variables in the abstract program that are not related to each other). The worst-case, however, is to reach full elaboration.

Then, an SCT decision procedure has to be applied. What SCT procedure? It is interesting to consider the closure-based algorithm [26], the SCT version of Algorithm 4.13;

this is apparently the most popular SCT algorithm (as witnessed by, e.g., [37, 30, 18]). Let \mathcal{A} be the stabilized system. Suppose that it has m flow-points. Its closure set can reach (at the worst case) a size of approximately $m^2 3^{n^2}$, as there are 3^{n^2} possible size-change graphs for a given pair of flow-points. The complexity of the algorithm is thus bounded by $m^2 3^{n^2} \leq (B_n)^2 3^{n^2}$ times some low-order polynomial, so, succinctly, it is $2^{\Theta(n^2)}$.

Surprisingly, the worst-case complexity *drops* if we assume that the MCS has been fully elaborated (at first sight the wasteful choice). In fact, fixing a pair of flow-points, there are less than n^{2n} different graphs between them that have the downward closure property (Page 18). This property is preserved by composition, so the closure computed in deciding SCT will be bounded by $m^2 n^{2n}$, where $m \leq B_n = 2n^{n-1}$, so the total complexity will be $2^{O(n \log n)}$. Thus, full elaboration improves the asymptotic worst-case complexity of the closure algorithm. However, as the next section will show, if we have a fully-elaborated system, there is actually a polynomial-time algorithm that decides termination (and more); this is because the SCT instance obtained is of a special structure. Hence, there is no need to do anything as costly as a closure computation.

5.7. The Codish-Lagoon-Stuckey Algorithm. The decision algorithm suggested by [12] is to first compute the closure set of a given MCS, and then, to every cyclic MC, apply “balancing” which is similar to our stabilization but on a local basis. After balancing, the cyclic graphs are tested, as for SCT. Thus, the algorithm is very similar to Algorithm 4.13; the LTT is replaced by the balancing procedure followed by Sagiv’s test (which as in SCT takes a simpler form for idempotent graphs).

Clearly, this test must be equivalent to the LTT in its results; its complexity may differ, though *grosso modo* they are similar—both are low-order polynomial in the size of the MC.

Codish, Lagoon and Stuckey do not prove completeness of their analysis with respect to termination of the original system, but to termination of each tested graph as a singleton; however this gap is not hard to bridge. Thus, we have at our disposal three slightly different algorithms for deciding MCS termination (four, when the results of the next section are taken into account); but it may suffice to take only two home: one closure-based algorithm (I propose Algorithm 4.13), and one based on full-elaboration (and the continuation to be given in the next section). The first has a higher worst-case complexity, $m^2 2^{\Theta(n^2)}$ versus $m^2 2^{\Theta(n \log n)}$; but it has a lower best-case complexity, which may be useful in practice. Moreover, its upper bound drops to $m^2 2^{\Theta(n \log n)}$ for a useful class of SCT instances, fan-in free graphs [18].

A clarification may be due regarding the assumptions made on *Val*. Section 2.2 includes the statement “all notions and results in this paper work equally well with partial orders, and even partial quasi-orders.” Full elaboration, as well as the “economic” stabilization, assume that the order is total; for any $u, v \in Val$, one of $x < y$, $x = y$ and $x > y$ must hold. However, by Lemma 3.6, it is possible to extend a given partial order to a total one. While this extension is not necessarily constructive, this does not matter: it suffices to imagine that such an extension is being used in order to explain the algorithm. After all, the termination condition does not really depend on the semantic domain, since it can be stated in pure graph terms (Section 3.1).

The results of the following section are an exception: if we ask, not only for decision regarding termination, but for a ranking function, the order on *Val* must really be total, since the ranking function descends in an order derived from it.

6. CONSTRUCTING A GLOBAL RANKING FUNCTION

This section describes a ranking-function construction for monotonicity constraint systems. At the heart of the construction is an algorithm that constructs a ranking function, if possible, for a fully-elaborated system. It does so in polynomial time, which highlights the fact that fully-elaborated systems are a special class of SCT instances. They are also special in that their global ranking functions are particularly simple—a tuple of elements (constants or variables) is associated with each flow-point, and these descend lexicographically in each and every transition.

The rest of this section is organized as follows. The first subsection presents tools that are used in exposing the structure of a fully-elaborated system. The second employs them to construct ranking functions. The third puts together the result for general monotonicity constraint systems.

6.1. Thread preservers and freezers. In preparation for the analysis of fully elaborated systems, we define complete threads, thread preservers and freezers and prove some essential properties. Note: when notions of connectivity are applied to an MCS (for example, “ \mathcal{A} is strongly connected”), they actually concern the underlying control-flow graph.

Definition 6.1 (complete thread). A thread in a given multipath is *complete* if it starts at the beginning of the multipath, and is as long as the multipath.

Lemma 6.2. *If a strongly connected MCS satisfies SCT, every finite multipath includes a complete thread.*

Proof. Assume that SCT is satisfied by the given MCS; then every infinite multipath contains an infinite thread. Let M be any finite multipath, beginning and ending at the same flow-point. Consider M^ω , i.e., the concatenation of infinitely many copies of M . From an infinite thread in M^ω one can clearly “cut out” a complete thread in one of the copies of M . Thus, every finite multipath that begins and ends at the same flow-point has a complete thread. A multipath M' that does not end at the point of departure can be extended to a multipath that does, since we assumed strong connectivity of the MCS. Thus, it has a complete thread. \square

Definition 6.3 (thread preserver). Given MCS \mathcal{A} , a mapping $P : F^{\mathcal{A}} \rightarrow \mathcal{P}(\{1, \dots, n\})$ is called a *thread preserver* of \mathcal{A} if for every $G : f \rightarrow g$ in \mathcal{A} , it holds that whenever $i \in P(f)$, there is $j \in P(g)$ such that $x_i \rightarrow x'_j \in G$.

It is easy to see that the set of thread preservers of \mathcal{A} is closed under union. Hence, there is a unique maximal thread preserver, which we denote by $\text{MTP}(\mathcal{A})$. Given a standard representation of \mathcal{A} , $\text{MTP}(\mathcal{A})$ can be found in linear time (for details see [7]).

Definition 6.4. A variable x_i is called *thread-safe* at flow-point f if every finite \mathcal{A} -multipath, starting at f , includes a complete thread from x_i .

Lemma 6.5. *Let \mathcal{A} be a fully-elaborated, strongly connected, terminating constraint system. For every f , let $S(f)$ be the set of indices of variables that are thread-safe at f . Then $S(f)$ is not empty for any $f \in F^{\mathcal{A}}$ and S is a thread preserver.*

Proof. Let M be any finite \mathcal{A} -multipath starting at f . Observe that since \mathcal{A} satisfies SCT and is strongly connected, there must be a complete thread in M , say starting at x_i . But

then x_n can also start a thread (note the downward-closure of fully-elaborated MCs). It follows that $n \in S(f)$.

We now aim to show that S is a thread preserver. Let $i \in S(f)$, and let $G : f \rightarrow g$. Every finite multipath M beginning with G has a complete thread that begins with an arc from x_i , say $x_i \rightarrow x'_{j_M}$. Let J be the set of all such indices j_M , and $k = \max J$. Then $x_i \rightarrow x'_k$ is an arc of G , because $k \in J$; and by the downward-closure property one can see that every M has a complete thread beginning with the arc $x_i \rightarrow x'_k$. Hence, $k \in S(g)$ and the proof is complete. \square

Informally, with a fully elaborated system, we are assured by this lemma that non-empty thread preservers exist. The next lemma shows that given any thread preserver we can find, within it, a *singleton thread preserver*. This comes very close to identifying a ranking function (or at least a quasi-ranking function, defined in the next subsection).

Lemma 6.6. *Let \mathcal{A} be a fully-elaborated, terminating MCS, and S a thread preserver, where $S(f) \neq \emptyset$ for all f . For every $f \in F^{\mathcal{A}}$, let $i_f = \min S(f)$. Then $P(f) = \{i_f\}$ is a thread preserver. In other words, every MC $G : f \rightarrow g$ includes $x_{i_f} \rightarrow x'_{i_g}$.*

Proof. By the definition of a thread-preserver, G must have an arc from $x_{i_f} \rightarrow x'_j$ with $j \in S(g)$; so by downward-closure, G includes $x_{i_f} \rightarrow x'_{i_g}$. \square

Definition 6.7 (freezer). Let $C : F^{\mathcal{A}} \rightarrow \{1, \dots, n\}$ denote a choice of one variable for each flow-point. Such C is called a *freezer* for \mathcal{A} if for every $G \in \mathcal{A}$, $G \vdash x_{C(f)} = x'_{C(g)}$.

Informally, a freezer is a singleton thread-preserver where the values are “frozen” since all the arcs represent equality.

Lemma 6.8. *Let \mathcal{A} be a stable MCS that satisfies SCT, and has a freezer C . If for every f , variable $x_{C(f)}$ is ignored, SCT will still be satisfied.*

Proof. Let M be an infinite multipath of \mathcal{A} ; by the SCT property, M has an infinitely descending thread τ . Observe that C induces an infinite thread ϑ in M , consisting entirely of no-change arcs. We claim that ϑ can have at most finitely many intersections with τ . To prove it, assume the contrary. Then there must be a strict arc in τ between two intersections. That is, the nodes at the intersections are connected by a descending path (via τ) and by a no-change path (via θ), which is a contradiction, making this part of M unsatisfiable. But, according to Lemma 5.11, this cannot happen. It follows that M has an infinitely descending thread (namely τ , minus some finite prefix) that avoids the frozen variables. Since this holds for any infinite multipath, SCT is satisfied with these variables omitted. \square

6.2. Ranking functions for fully-elaborated systems. From this point on, fix \mathcal{A} to be a fully-elaborated system.

Provided that it terminates, a ranking function will be constructed. To precisely specify the form of this function, we define vectors (definition from [8], but simplified).

Definition 6.9 (vectors). Let $n > 0$ represent the number of variables in a program under consideration, and let $B > 0$ be some integer. V^B is the set of tuples $\mathbf{v} = \langle v_1, v_2, \dots \rangle$ of even length, where every even position is a variable among $\{x_1, \dots, x_n\}$, such that every variable appears at most once; and every odd position is an integer between 0 and B .

Definition 6.10. The value of $\mathbf{v} \in V^B$ in program state (f, σ) , denoted $\mathbf{v}\sigma$, is obtained by substituting the values of variables according to σ . This results in a tuple with elements of Val and integers in even and odd positions, respectively. Such tuples are compared lexicographically.

The functions we construct have a simple form: to each flow-point f a vector \mathbf{v}_f is associated so that $\rho(f, \sigma) = \mathbf{v}_f \sigma$. Thus, for every transition $(f, \sigma) \mapsto (g, \sigma')$, we shall have $\mathbf{v}_f \sigma \succ \mathbf{v}_g \sigma'$, where \succ is the strict lexicographic order.

The construction is incremental. To justify the incremental construction, we define quasi ranking functions and residual transition systems.

Definition 6.11. Let \mathcal{T} be a transition system with state space St . A *quasi-ranking function* for \mathcal{T} is a function $\rho : St \rightarrow W$, where W is a well-founded set, such that $\rho(s) \geq \rho(s')$ for every $(s, s') \in \mathcal{T}$.

The residual transition system relative to ρ , denoted \mathcal{T}/ρ , includes all (and only) the transitions of \mathcal{T} which do *not* decrease ρ .

Note that when Lemma 6.6 applies, it provides a quasi-ranking function: $\rho(f, \sigma) = \langle x_{i_f} \rangle \sigma$.

The next couple of lemmas are quite trivial but we spell them out because they clarify how a ranking function may be constructed incrementally. We consider the codomain of all our functions to consist of lexicographically-ordered tuples over “scalars” (a scalar is either a constant or a variable) and we use the notation $v ++ u$ for concatenation of tuples.

Lemma 6.12. Assume that ρ is a quasi-ranking function for \mathcal{T} , and ρ' a ranking function for \mathcal{T}/ρ ; then $\rho ++ \rho'$ is a ranking function for \mathcal{T} .

Lemma 6.13. Assume that the CFG of \mathcal{A} consists of a set C_1, \dots, C_k of mutually disconnected components (that is, there is no arc from C_i to C_j with $i \neq j$). If for every i , ρ_i is a ranking function for \mathcal{A} restricted to C_i , then $\cup_i \rho_i$ is a ranking function for \mathcal{A} .

Lemma 6.14. Suppose that the CFG of \mathcal{A} consists of several strongly connected components (SCCs). Let C_1, \dots, C_k be a reverse topological ordering of the components. Define a function $\rho(s)$ for $s = (f, \sigma)$ as the index i of the component C_i including f . Then ρ is a quasi-ranking function (with co-domain $[1, k]$) and it is strictly decreasing on every transition represented by an inter-component arc.

The following algorithm puts all of this together. Note: a CFG whose arc set is empty is called *vacant*. A strongly connected component whose arc set is empty is called *trivial* (it may have connections to other components).

Algorithm 6.15. (ranking function construction for \mathcal{A})

The algorithm assumes that \mathcal{A} is fully elaborated. If \mathcal{A} terminates, a ranking function will be returned. Otherwise, the algorithm will fail.

We assume that the representation of \mathcal{A} allows for “hiding” certain variables of any given flow point. This affects subsequent MTP computations, which will ignore the hidden variables.

- (1) List the SCCs of \mathcal{A} in reverse-topological order. For each $f \in F^{\mathcal{A}}$, let κ_f be the position of the SCC of f . Form \mathcal{A}' by deleting all the inter-component transitions. If \mathcal{A}' is vacant, return ρ where $\rho(f, \sigma) = \kappa_f$.

- (2) For each SCC C , compute the MTP, using the algorithm in [7]. If empty, report failure and exit.
- (3) For every f , let x_{i_f} be the lowest MTP variable of f .
- (4) For every graph $G : f \rightarrow g$, if it includes $x_{i_f} \xrightarrow{>} x'_{i_g}$, delete the graph from \mathcal{A}' ; otherwise, retain the graph but hide x_{i_f} .
- (5) For every f , let $\rho(f, \sigma) = \langle \kappa_f, x_{i_f} \rangle \sigma$.
- (6) If \mathcal{A}' is now vacant, return ρ . Otherwise, compute a ranking function ρ' recursively for \mathcal{A}' , and return $\rho ++ \rho'$.

Correctness. We claim that the abstract program \mathcal{A}' , passed to the recursive call, always represents the residual transition system $\mathcal{T}_{\mathcal{A}}/\rho$. This should be clear when we delete graphs that strictly decrease ρ . The less-trivial point is the treatment of graphs G where the MTP arc $x_{i_f} \rightarrow x_{i_g}$ is non-strict (Step 4). To obtain the residual system precisely we should have replaced the inequality constraints with equalities: $x_{i_f} = x'_{i_g}$. However, having done so, the set of indices $C(f) = i_f$ becomes a freezer, and therefore can be ignored (Lemma 6.8).

Hiding the “frozen” variables ensures that these variables will not be used again in ρ' . So in the final tuple $\rho(f, \sigma) ++ \rho'(f, \sigma)$, each variable will occur at most once. This shows, in particular, that the recursion always terminates, which means that the residual transition system is eventually vacant, and when this happens, we have a ranking function.

Complexity. The algorithm will make at most n passes in which two elements are added to the tuple. The costly part of a pass is the MTP computation which takes time linear in the sum of sizes of all the MC graphs, that is, $|\mathcal{A}|n^2$ (recall that $|\mathcal{A}|$ is the number of MCs). In subsequent passes, the system is diminished, so an upper bound on the total running time is $O(|\mathcal{A}|n^3)$.

6.3. Ranking functions for all! So far, we have constructed ranking functions for fully-elaborated systems. To construct a ranking function for a general MCS \mathcal{B} , we first transform it into a fully-elaborated \mathcal{A} using Algorithm 5.6. Then, Algorithm 6.15 can be applied. The ranking function for \mathcal{A} can be translated to one for \mathcal{B} , as the next lemma shows.

Lemma 6.16. *If \mathcal{A} simulates \mathcal{B} deterministically, any ranking function for \mathcal{A} can be transformed into a ranking function for \mathcal{B} .*

Proof. Let ρ be the \mathcal{A} ranking function. The \mathcal{B} ranking function is defined for state (f, σ) as $\rho(g, \sigma')$ where g is the unique point such that $(f, g) \in \phi$ and $\sigma' \models I_g$, where $\sigma' = \sigma \circ (\psi_{g,f})$. \square

When \mathcal{A} is fully elaborated, we have a function ρ where $\rho(g, \sigma)$ is a fixed vector \mathbf{v}_g for every $g \in F^{\mathcal{A}}$; thus the resulting \mathcal{B} ranking function has the form

$$\rho(f, \sigma) = \begin{cases} \mathbf{v}_{g_1}(\sigma \circ \psi_{g,f}) & \text{if } \sigma \circ \psi_{g,f} \models I_{g_1} \\ \mathbf{v}_{g_2}(\sigma \circ \psi_{g,f}) & \text{if } \sigma \circ \psi_{g,f} \models I_{g_2} \\ \dots & \\ \mathbf{v}_{g_t}(\sigma \circ \psi_{g,f}) & \text{if } \sigma \circ \psi_{g,f} \models I_{g_t} \end{cases}$$

where g_1, \dots, g_t are the \mathcal{A} flow-points associated with \mathcal{B} flow-point f .

This function may be simplified by combining rows that have the same vector as the function’s value, so we do not have to list B_n rows. One example of such a function appears




Figure 6: MCs for ranking function example. There is a single flow-point.

in the introduction (conclusion of Example 2.3). For another example, consider the MCs depicted in Figure 6. A ranking function of the form we consider is

$$\rho(x_1, x_2, x_3, x_4) = \begin{cases} \langle 1, x_1, 1, x_3 \rangle & \text{if } x_1 > x_2 \\ \langle 1, x_2, 1, x_4 \rangle & \text{if } x_1 \leq x_2. \end{cases}$$

(obviously, the 1-valued entries can be dropped).

Theorem 6.17. *Let \mathcal{B} be a terminating MCS, with m flow-points and n variables per point. There is a ranking function ρ for \mathcal{B} where $\rho(f, \sigma)$ is defined by a set of elements of V^{mB_n} , each one associated with certain inequalities on variables, which define the region where the given vector is the function value. There are at most B_n different vectors for any flow-point. The complexity of constructing ρ is $O(|\mathcal{B}| \cdot n^{2n+1})$.*

Proof. Fully-elaborating \mathcal{B} takes $O(|\mathcal{B}|n^{2n+1})$ time and creates a system with $O(|\mathcal{B}|n^{2n-2})$ MCs. The running time of Algorithm 6.15 on the result is $O(|\mathcal{A}|n^3) = O(|\mathcal{B}|n^{2n+1})$. \square

Remarks. This is the first algorithm to construct explicit ranking functions for any terminating MCS, but even when restricting attention to SCT instances, the results improve upon previous publications. The improvement over [8] is that any positive instance can be handled; the improvement over [25] is that in that work, the vectors were possibly doubly exponential in length (as a function of n) and the complexity of the construction was only bounded by a triply exponential function.

The complexity of our construction is optimal (or very nearly so) in two senses.

- [8] considers ranking functions of the form described above, and shows that there are SCT instances with n variables that require $(n-1)!$ distinct vectors per flow-point. Our construction yields the close upper bound B_n .
- In a vector, each variable appears at most once. This is clearly optimal since they may all be involved in the termination proof. This also means that the codomain of the ranking functions has the smallest necessary ordinal (in general).

Finally let us remark that since the processing of a fully-elaborated system is polynomial-time, it is advantageous to use it rather than applying any general SCT decision procedure, even if all we want is a yes-no answer.

7. ROOTED VERSUS UNIFORM TERMINATION

In this paper, the notion of termination used was *uniform termination*, which means that there must be no cycles in the whole state space of the modelled transition system. Practically, it may be desirable to account for *rooted termination*, when only computation

paths beginning at a given initial point f_0 (and satisfying its state invariant) are considered, thus avoiding some false alarms.

In this paper, uniform termination was chosen in favor of simplicity. However, It is not be difficult to adapt our methods to rooted termination; [26] shows how to do that for SCT.

8. CONCLUSION AND RESEARCH QUESTIONS

We studied the MCS abstraction, an appealing extension of the Size-Change Termination framework, showing how several elements of the theory of SCT can be achieved in the stronger framework: sound and complete termination criteria, closure-based algorithms based on local termination tests, and the construction of global ranking functions (which improves on previously-published work even for SCT). A key technique was refining the abstraction, using state invariants to separate out the behaviour under different assumptions on the relative order of variables. We showed that MCS termination can be simplified in this way not only to SCT termination, but to an easy case thereof.

The contribution of this paper is theoretical; hopefully, it will trigger further research, moving towards the practical application of the theory.

The algorithms in this article were aimed at simplicity of presentation and analysis and can certainly be improved. If the purpose is just to decide termination (and possibly produce a non-terminating multipath if one exists), there is a choice between a closure-based algorithm (Section 4) and an algorithm based on full elaboration (Section 6). While theoretically, full elaboration reduces the exponent in the worst-case time from $\Theta(n^2)$ to $\Theta(n \log n)$, avoiding it may well be more efficient in practice. The intriguing effectiveness of the (theoretically worst) closure-based algorithms was discussed, in the context of SCT, by Fogarty and Vardi [18].

Algorithm 6.15 has been implemented in Java, but so far has only been tried out with toy examples, so it is too early to speak of an empirical evaluation. As expected, memory fills up quickly when the number of variables is enlarged. It seems clear that in a practical implementation, both for deciding termination and for constructing ranking functions, avoiding unnecessary combinatorial explosion is imperative. Some tactics that should probably be used include an initial analysis to identify the subset of variables that are pertinent to termination [3, 29], and—obviously—handling one SCC of the original system at a time, instead of fully-elaborating all at once. Furthermore, we may choose to resort to heuristics that sacrifice completeness for efficiency. In the context of SCT analysis, some heuristics have been studied [11, 7, 6], and similar approaches can conceivably be useful with MCS.

A possible extension of this work is modeling arbitrary Boolean combinations of order constraints (including disjunctions) in a direct way (rather than by converting to disjunctive normal form, as discussed in Section 2.1). It may be interesting to find out if this has practical interest.

Finally, perhaps the most appealing aspect of MCS, compared to the SCT abstraction, is its usefulness in the integer domain. Adapting the theory of monotonicity constraint termination to the integer domain is to be the subject of a forthcoming paper.

Acknowledgments. The author thanks Chin Soon Lee for inspiration; the LIPN laboratory at Université Paris 13, where part of this work was done, for hospitality; Ariel Snir for comments on the manuscript (and for programming); the editor, Neil Jones, who pointed out

an error in a previous version, and the anonymous referees, who have provided exceptionally thorough reviews and many suggestions for improvement.

REFERENCES

- [1] Elvira Albert, Puri Arenas, Michael Codish, Samir Genaim, Germán Puebla, and Damiano Zanardini. Termination analysis of java bytecode. In Gilles Barthe and Frank S. de Boer, editors, *Formal Methods for Open Object-Based Distributed Systems, 10th IFIP WG 6.1 International Conference, FMOODS 2008, Oslo, Norway, Proceedings*, volume 5051 of *Lecture Notes in Computer Science*, pages 2–18. Springer, 2008.
- [2] Elvira Albert, Puri Arenas, Samir Genaim, and Germán Puebla. Automatic inference of upper bounds for recurrence relations in cost analysis. In María Alpuente and Germán Vidal, editors, *Static Analysis, 15th International Symposium, SAS 2008, Valencia, Spain, Proceedings*, volume 5079 of *Lecture Notes in Computer Science*, pages 221–237. Springer, 2008.
- [3] Elvira Albert, Puri Arenas, Samir Genaim, Germán Puebla, and Damiano Zanardini. Removing useless variables in cost analysis of java bytecode. In *SAC*, pages 368–375, 2008.
- [4] James Avery. Size-change termination and bound analysis. In M. Hagiya and P. Wadler, editors, *Functional and Logic Programming: 8th International Symposium, FLOPS 2006*, volume 3945 of *Lecture Notes in Computer Science*. Springer, 2006.
- [5] Amir M. Ben-Amram. Size-change termination, monotonicity constraints and ranking functions. In Ahmed Bouajjani and Oded Maler, editors, *Computer Aided Verification, 20th International Conference, CAV 2009, Grenoble, France, volume 5643 of Lecture Notes in Computer Science*, pages 109–123. Springer, 2009.
- [6] Amir M. Ben-Amram and Michael Codish. A SAT-based approach to size change termination with global ranking functions. In C.R. Ramakrishnan and Jakob Rehof, editors, *14th Intl. Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 5028 of *LNCS*, pages 46–55. Springer, 2008.
- [7] Amir M. Ben-Amram and Chin Soon Lee. Size-change analysis in polynomial time. *ACM Transactions on Programming Languages and Systems*, 29(1), -*ACM Transactions on Programming Languages and Systems*, 29(1), 2007.
- [8] Amir M. Ben-Amram and Chin Soon Lee. Ranking functions for size-change termination II. *Logical Methods in Computer Science*, 5(2), 2009.
- [9] Florence Benoy and Andy King. Inferring argument size relationships with CLP(R). In John Gallagher, editor, *Logic Program Synthesis and Transformation (Proceedings)*, volume 1207 of *Lecture Notes in Computer Science*, pages 204–223. Springer-Verlag, August 1997.
- [10] Alexander Brodsky and Yehoshua Sagiv. Inference of inequality constraints in logic programs. In *Proceedings of the Tenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS), 1991*, pages 227–240. ACM Press, 1991.
- [11] Michael Codish, Vitaly Lagoon, Peter Schachte, and Peter J. Stuckey. Size-change termination analysis in k -bits. In Peter Sestoft, editor, *Programming Languages and Systems, 15th European Symposium on Programming, ESOP 2006, Vienna, Austria, Proceedings*, volume 3924 of *Lecture Notes in Computer Science*, pages 230–245. Springer, 2006.
- [12] Michael Codish, Vitaly Lagoon, and Peter J. Stuckey. Testing for termination with monotonicity constraints. In Maurizio Gabbriellini and Gopal Gupta, editors, *Logic Programming, 21st International Conference, ICLP 2005*, volume 3668 of *Lecture Notes in Computer Science*, pages 326–340. Springer, 2005.
- [13] Michael Codish and Cohavit Taboch. A semantic basis for termination analysis of logic programs. *The Journal of Logic Programming*, 41(1):103–123, 1999. preliminary (conference) version in LNCS 1298 (1997).
- [14] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, Cambridge, MA, 2001.
- [15] Patrick Cousot and Nicholas Halbwachs. Automatic discovery of linear restraints among variables of a program. In *Conference Record of the Fifth annual ACM Symposium on Principles of Programming Languages*, pages 84–96. ACM, ACM, January 1978.
- [16] Danny de Schreye and Stefaan Decorte. Termination of logic programs: The never-ending story. *Journal of Logic Programming*, 19-20:199–260, May-July 1994.

- [17] Nachum Dershowitz, Naomi Lindenstrauss, Yehoshua Sagiv, and Alexander Serebrenik. A general framework for automatic termination analysis of logic programs. *Applicable Algebra in Engineering, Communication and Computing*, 12(1–2):117–156, 2001.
- [18] Seth Fogarty and Moshe Y. Vardi. Büchi complementation and Size-Change Termination. In *Proceedings of the 15th Intl. Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 5505 of *Lecture Notes in Computer Science*, pages 16–30. Springer, 2009.
- [19] Andrew V. Goldberg. Scaling algorithms for the shortest paths problem. *SIAM Journal on Computing*, 24(3):494–504, June 1995.
- [20] Neil D. Jones. Automatic program specialization: A re-examination from basic principles. In D. Björner, A. P. Ershov, and N. D. Jones, editors, *Partial evaluation and Mixed Computation*, pages 225–282, 1988.
- [21] Neil D. Jones, Carsten K. Gomard, and Peter Sestoft. *Partial Evaluation and Automatic Program Generation*. Prentice Hall, 1993.
- [22] Alexander Krauss. Certified size-change termination. In Frank Pfenning, editor, *11th International Conference on Automated Deduction (CADE)*, volume 4603 of *LNAI*, pages 460–475. Springer-Verlag, July 2007.
- [23] Alexander Krauss and Armin Heller. A mechanized proof reconstruction for SCNP termination. Presented in the Tenth International Workshop on Termination WST’09, Leipzig, 2009.
- [24] Gary T. Leavens, Albert L. Baker, and Clyde Ruby. JML: A notation for detailed design. In Haim Kilov, Bernhard Rumpe, and Ian Simmonds, editors, *Behavioral Specifications of Businesses and Systems*, pages 175–188. Kluwer Academic Publishers, Boston, 1999.
- [25] Chin Soon Lee. Ranking functions for size-change termination. *ACM Transactions on Programming Languages and Systems*, 31(3), April 2009.
- [26] Chin Soon Lee, Neil D. Jones, and Amir M. Ben-Amram. The size-change principle for program termination. In *Proceedings of the Twenty-Eighth ACM Symposium on Principles of Programming Languages, January 2001*, volume 28, pages 81–92. ACM press, January 2001.
- [27] Naomi Lindenstrauss and Yehoshua Sagiv. Automatic termination analysis of Prolog programs. In Lee Naish, editor, *Proceedings of the Fourteenth International Conference on Logic Programming*, pages 64–77, Leuven, Belgium, Jul 1997. MIT Press.
- [28] Naomi Lindenstrauss, Yehoshua Sagiv, and Alexander Serebrenik. Proving termination for logic programs by the query-mapping pairs approach. In Maurice Bruynooghe and Kung-Kiu Lau, editors, *Program Development in Computational Logic: A Decade of Research Advances in Logic-Based Program Development*, volume 3049 of *Lecture Notes in Computer Science*, pages 453–498. Springer, 2004.
- [29] Panagiotis Manolios and Aaron Turon. All-Termination(T). In *Proceedings of the 15th Intl. Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 5505 of *Lecture Notes in Computer Science*, pages 398–412. Springer, 2009.
- [30] Panagiotis Manolios and Daron Vroon. Integrating static analysis and general-purpose theorem proving for termination analysis. In *Proceedings of ICSE’06: International Conference on Software Engineering*, pages 873–876. ACM, 2006.
- [31] Panagiotis Manolios and Daron Vroon. Termination analysis with calling context graphs. In *Proceedings, Computer Aided Verification, 18th International Conference, CAV 2006, Seattle, WA, USA*, volume 4144 of *LNCIS*, pages 401–414. Springer-Verlag, 2006.
- [32] Lee Naish. Automating control for logic programs. *Journal of Logic Programming*, 2(3):167–183, October 1985.
- [33] Andreas Podelski and Andrey Rybalchenko. Transition invariants. In Harald Ganzinger, editor, *LICS’04: Logic in Computer Science*, pages 32–41. IEEE Computer Society, 2004.
- [34] Yehoshua Sagiv. A termination test for logic programs. In Vijay Saraswat and Kazunori Ueda, editors, *Logic Programming, Proceedings of the 1991 International Symposium, San Diego, California, USA*, pages 518–532. MIT Press, 1991.
- [35] N. J. A. Sloane. The on-line encyclopedia of integer sequences. Published electronically at www.research.att.com/~njas/sequences/.
- [36] Fausto Spoto, Fred Mesnard, and Étienne Payet. A termination analyzer for Java bytecode based on path-length. *ACM Trans. Program. Lang. Syst.*, 32(3):1–70, 2010.
- [37] René Thiemann and Jürgen Giesl. The size-change principle and dependency pairs for termination of term rewriting. *Applicable Algebra in Engineering, Communication and Computing*, 16(4):229–270, September 2005.